

**AFRL-IF-RS-TR-2004-63**  
**Final Technical Report**  
**March 2004**



# **TRANSITIONING SECURE BORDER GATEWAY PROTOCOL (S-BGP) INTO THE INTERNET**

**BBN Technologies**

**Sponsored by**  
**Defense Advanced Research Projects Agency**  
**DARPA Order No. K496**

*APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.*

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

**AIR FORCE RESEARCH LABORATORY**  
**INFORMATION DIRECTORATE**  
**ROME RESEARCH SITE**  
**ROME, NEW YORK**

## **STINFO FINAL REPORT**

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2004-63 has been reviewed and is approved for publication

APPROVED:       /s/

ROBERT L. KAMINSKI  
Project Engineer

FOR THE DIRECTOR:       /s/

WARREN H. DEBANY, JR., Technical Advisor  
Information Grid Division  
Information Directorate

<b>REPORT DOCUMENTATION PAGE</b>			Form Approved OMB No. 074-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503				
<b>1. AGENCY USE ONLY (Leave blank)</b>		<b>2. REPORT DATE</b> MARCH 2004	<b>3. REPORT TYPE AND DATES COVERED</b> Final Aug 00 – Jan 04	
<b>4. TITLE AND SUBTITLE</b> TRANSITIONING SECURE BORDER GATEWAY PROTOCOL (S-BGP) INTO THE INTERNET			<b>5. FUNDING NUMBERS</b> C - F30602-00-C-0212 PE - 62301E PR - K496 TA - 91 WU - A1	
<b>6. AUTHOR(S)</b> Stephen T. Kent, Charles W. Lynn and Karen S. Seo				
<b>7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)</b> BBN Technologies 10 Moulton Street Cambridge Massachusetts 02138			<b>8. PERFORMING ORGANIZATION REPORT NUMBER</b>  N/A	
<b>9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)</b> Defense Advanced Research Projects Agency AFRL/IFG 3701 North Fairfax Drive 525 Brooks Road Arlington Virginia 22203-1714 Rome New York 13441-4505			<b>10. SPONSORING / MONITORING AGENCY REPORT NUMBER</b>  AFRL-IF-RS-TR-2004-63	
<b>11. SUPPLEMENTARY NOTES</b>  AFRL Project Engineer: Robert L. Kaminski/IFG/(315) 330-1815/ Robert.Kaminski@rl.af.mil				
<b>12a. DISTRIBUTION / AVAILABILITY STATEMENT</b> APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.				<b>12b. DISTRIBUTION CODE</b>
<b>13. ABSTRACT (Maximum 200 Words)</b> Internet routing is based on a distributed system composed of many routers grouped into management domains called Autonomous Systems (ASes). Routing information is exchanged between ASes in Border Gateway Protocol (BGP) UPDATE messages. BGP is a critical component of the Internet's routing infrastructure. However, it is highly vulnerable to a variety of attacks due to the lack of a scalable means of verifying the authenticity and authorization of BGP control traffic. Secure BGP (S-BGP) addresses these vulnerabilities. The S-BGP architecture employs three security mechanisms. First, a Public Key Infrastructure (PKI) is used to support the authentication of ownership of IP address blocks, ownership of Autonomous System (AS) numbers, and a BGP router's identity and its authorization to represent as AS. Second, a new, optional, GBP transitive path attribute is employed to carry digital signatures ("route attestations") covering the routing information in a BGP UPDATE. Third, IPsec is used to provide data and partial sequence integrity, and to enable BGP routers to authenticate each other for exchanges of BGP control traffic.				
<b>14. SUBJECT TERMS</b> Border Gateway Protocol, Network Security, Network Routing, Critical Infrastructure Protection			<b>15. NUMBER OF PAGES</b> 16	
			<b>16. PRICE CODE</b>	
<b>17. SECURITY CLASSIFICATION OF REPORT</b>  UNCLASSIFIED	<b>18. SECURITY CLASSIFICATION OF THIS PAGE</b>  UNCLASSIFIED	<b>19. SECURITY CLASSIFICATION OF ABSTRACT</b>  UNCLASSIFIED	<b>20. LIMITATION OF ABSTRACT</b>  UL	

## Table of Contents

1	Summary of technical objectives and approach.....	1
1.1	Overview and objectives .....	1
1.2	Approach.....	1
2	Accomplishments .....	2
2.1	S-BGP Software .....	2
2.2	Work with Regional Internet Registries (RIRs).....	4
2.3	Work with ISPs .....	4
2.4	Workshops .....	5
2.5	Papers, Internet Drafts, Reports .....	6
2.6	Meetings and Briefings .....	7
3	Papers, Presentations, Specifications, and Code.....	8
3.1	Papers .....	8
3.2	Presentations .....	8
3.4	Specifications .....	8
3.5	Source Code .....	8
4	Lessons Learned.....	9
5	BGP Statistics Update .....	11

# **1 Summary of technical objectives and approach**

## **1.1 Overview and objectives**

Internet routing is based on a distributed system composed of many routers, grouped into management domains called Autonomous Systems (ASes). Routing information is exchanged between ASes in Border Gateway Protocol (BGP) UPDATE messages. BGP is a critical component of the Internet's routing infrastructure. However, it is highly vulnerable to a variety of attacks due to the lack of a scalable means of verifying the authenticity and authorization of BGP control traffic. Secure BGP (S-BGP) addresses these vulnerabilities.

The S-BGP architecture employs three security mechanisms. First, a Public Key Infrastructure (PKI) is used to support the authentication of ownership of IP address blocks, ownership of Autonomous System (AS) numbers, and a BGP router's identity and its authorization to represent an AS. This PKI parallels the IP address and AS number assignment system and takes advantage of the existing infrastructure (Regional Internet Registries, etc.) The PKI also supports digitally signed “address attestations” that permit address block owners to identify the AS(es) authorized to originate routes for these address block. Second, a new, optional, BGP transitive path attribute is employed to carry digital signatures (“route attestations”) covering the routing information in a BGP UPDATE. These signatures, along with certificates and address attestations from the S-BGP PKI, enable the receiver of a BGP routing UPDATE to verify the address prefixes and path information that it contains. Third, IPsec is used to provide data and partial sequence integrity, and to enable BGP routers to authenticate each other for exchanges of BGP control traffic.

Under a previous contract with DARPA, a proof-of-concept prototype of S-BGP was developed and used to demonstrate the effectiveness and feasibility of deploying S-BGP. However, a major obstacle to the deployment of S-BGP is that it requires the participation of several distinct organizations — the Internet registries, router vendors, and Internet service providers (ISPs). Because there will be no security benefits unless a few of each type of the organizations participate, each organization cannot justify the expense of investing in this new technology unless the others have also done so — a classic chicken-and-egg problem. The goal of this project is to overcome these obstacles and promote deployment of S-BGP into the Internet.

## **1.2 Approach**

Deploying S-BGP will require working with the Internet registries and ISPs to set up the PKI; working with router vendors to implement the S-BGP enhancements (new path attribute, IPsec, etc.) on COTS routers, and convincing ISPs to either buy and use these routers or to use an ancillary device (associated with a border router) to support S-BGP. To do this, BBN intends to take the following steps:

***Setting up the Public Key Infrastructure*** -- BBN will modify an existing certificate management system to support the X.509 v3 certificate extensions that S-BGP uses as a basis for authorization and to enforce the S-BGP hierarchical address and AS number delegation constraints before signing a subordinate certificate. This system is being used as the basis for the open source certificate management system being developed for the CHATS program. The S-BGP enhancements will be incorporated into the CHATS CMS software and the resulting system will be freely available to internet registries, ISPs, DSPs, etc. Additional tools/systems will be developed for the distribution of the resulting certificates.

***COTS implementation of S-BGP*** -- BBN will enhance the prototype S-BGP software to be more robust and to support features that were not needed in the proof-of-concept testing, e.g., multi-protocol support (IPv6) and communities. The availability of working code (a reference implementation) will reduce the cost of integrating S-BGP into routers. This software could also provide a code base for an outboard box which would run S-BGP and be operated in parallel with the existing router. BBN will also enhance the current S-BGP protocol specification to reflect "lessons learned" from implementation efforts, experience with ISPs, etc. As needed, BBN will provide guidance to router vendors on their implementation efforts and will develop a test suite to assess interoperability.

***ISP adoption of S-BGP*** -- BBN will use the existing S-BGP proof of concept prototype to create a test system, e.g., on a PC, that can be run in parallel with a real BGP speaker without interfering with the operational networks. In addition, tools will be developed to support the NOC operations that will be needed for S-BGP, e.g., downloading and validation of certificates and creation of certificate extracts to be pushed to the ISP's S-BGP routers.

## **2 Accomplishments**

### **2.1 S-BGP Software**

- ❖ *Reference Implementation and Test System* – We enhanced the prototype software to be more robust and added support for features such as route aggregation. We added S-BGP policy mechanisms and better support for testing and experimentation, e.g., more instrumentation. We also modified the system slightly to make it easier to demonstrate, e.g., made the routing table display clearer.
- ❖ *Certificate Management Tool Kit (CMTK) and CHATS CMS* -- The GUI and backend code of the CMTK were modified to support the three S-BGP private certificate extensions (IP addresses, AS numbers, Router ID). These changes were later integrated into the CHATS CMS.
- ❖ *Sign Assure Plus (formerly known as Super SafeKeyper) and the Rule Editor* -- The Sign Assure Plus was modified to extract the S-BGP extensions from a CA's certificate and use the values from the extension to verify that the IP Address ranges and AS number ranges in a certificate to be signed are subsets of the corresponding ranges in

the certificate of the signing CA. The “Rule Editor” used to create/edit the rules enforced by the SignAssure+ was enhanced to support rules that use the S-BGP certificate extensions. Enforcement of these rules by the SignAssure+ will ensure that a certification authority that has been assigned an IP address block or AS#s, can allocate addresses or AS #s ONLY from those it owns to subordinate organizations and subscribers.

- ❖ *Certificate/Address Attestation (AA) Server* (aka distributed repository system) -- An initial architecture and design were completed. An initial version of the database was completed. This system was ported from Windows to Linux and transaction upload/download processing was completed. An initial version of the certificate repository software was completed.
- ❖ *NOC Tools* -- We completed a release of the NOC tools that covers initial versions of the following functions:
  - Management of the tools
    - adding/deleting users, changing their passwords, and their privileges
    - installing trust anchors, etc.
  - Generation of a certificate request:
    - a CA certificate request that can be submitted with a request for IP address blocks or AS numbers, and subsequent storage of the resulting signed certificate in the local CA and in the local database.
    - certificate requests are supported for the following end-entity certificates:
      - Operator certificates (1 per Operator)
      - Network certificates (1 per Network)
      - Autonomous System (AS) certificates (1 per AS# being used by an Organization)
      - Router certificates (1 per S-BGP speaker)
      - IPsec certificates (1 per S-BGP speaker)
  - Generation and Signing of Address Attestations
  - Management of S-BGP certificates, CRLs, and AAs
    - adding locally-generated certificates, CRLs, and AAs to the local database
    - uploading local certificates, CRLs, and AAs to distributed repository – includes both adding and deleting certificates, CRLs, and AAs.
    - downloading certificates, CRLs, and AAs from distributed repository – The NOC tools include processing of the downloaded file in this step – validate, reconcile with local database, creation and signing of an extract file.
  - Updating of S-BGP countermeasures data in routers
    - uploading extract file – The NOC tools support a script for exporting extraction files to routers and also support a database of the organization’s routers indicating what extraction versions have been uploaded to which routers.

## **2.2 Work with Regional Internet Registries (RIRs)**

- ❖ On 10/17/00, Stephen Kent met with key staff at American Registry for Internet Numbers (ARIN) to discuss collaboration on setting up an initial PKI for S-BGP. ARIN staff present were:

- Ray Plzak (president and CEO)
- Leslie Nobile (registration services group manager)
- Richard Jimmerson (director of operations)
- Catherine Murphy (software engineer).

BBN later worked with ARIN to learn about their procedures and policies for processing requests for IP addresses and AS numbers and to define the S-BGP-related data fields they will need to add to their database for tracking IP address and AS number assignments. They provided a dump of their database (IP address allocations, AS number allocations, etc.) and a description of their procedures and policies for processing requests for IP addresses and AS numbers. In July 2001, we met with ARIN's president and senior staff to discuss

- issues related to initialization of the cert/AA database from ARIN's database dump
- how to integrate support for the S-BGP PKI with ARIN's procedures for assignment of IP addresses and AS numbers.

On 3/19/02, at Ray Plzak's (President of ARIN) invitation, Steve Kent briefed S-BGP to the ARIN Board of Directors. On 4/9/02, Steve Kent gave a presentation on S-BGP to ARIN user community.

- ❖ During 2002, we corresponded with representatives from APNIC (The Asia-Pacific RIR) about S-BGP. They are interested both in acting as a root in the S-BGP PKI and in the S-BGP certificate extensions (IP address and AS number) for more general use.
- ❖ In May 2003, Steve Kent gave a presentation on S-BGP to the RIPE (The European RIR, Réseaux IP Européens) community.<sup>1</sup>

## **2.3 Work with ISPs**

- ❖ We met/teleconferenced several times with staff from DISA to explore possible areas for collaboration on field experiments and to gain their input as an ISP on requirements for an S-BGP reference implementation and an outboard test system, e.g., routing policy. DISA participants included Keith Fuller, Lynn Keuthan, Nam Nguyen and David Coe. DISA agreed to collaborate on testing S-BGP and providing feedback on the associated NOC tools, however due to a facilities move, re-organization and other demands, they were not able to find time to do this work.

---

<sup>1</sup> This work was done under another contract but was relevant to this project and thus is mentioned here.



- ❖ Genuity agreed to evaluate/test S-BGP. On April 30, 2002, we delivered the first release of the reference/test implementation of S-BGP to Genuity Inc. Due to economic decline in the ISP arena, Genuity was unable to allocate the resources to test this system. Later, Genuity declared bankruptcy and its assets were acquired by Level3.
- ❖ In December, we briefed Peter Zarrella (advisor/deputy to Dawn Meyerriecks, the CTO/Chief Technologist of DISA), on S-BGP.

## **2.4 Workshops**

- ❖ Randy Bush (then IETF area director for Internet operations) arranged an opportunity for BBN to conduct an all-day, hands-on workshop on S-BGP to a half a dozen key engineers from influential service providers (October 30, 2002). The workshop was a success. The attendees wanted to assess whether or not S-BGP would really address BGP vulnerabilities and was deployable and scalable. BBN spent most of the day explaining various aspects of S-BGP, answering their questions, addressing "what if" scenarios, etc. In the end most of them went away thinking that S-BGP could be viable, which was a great improvement over their initial views.
- ❖ On January 17, 2003, BBN held a workshop on the Secure Border Gateway Protocol (S-BGP) in the DC area for representatives from various government agencies. The focus of this gathering was to provide background about S-BGP and answer questions about whether S-BGP really addresses BGP vulnerabilities and whether it is deployable and scalable given operational constraints. The presentation, demo, and subsequent discussions went very well. The following people attended:
  - Colonel Tim Gibson (US DoD, Joint Task Force -- Computer Network Operations)
  - Douglas Maughan (US DoD, Defense Advanced Research Projects Agency (DARPA))
  - Mike Ferguson (DARPA System Engineering and Technical Assistance (SETA) contractor)
  - Commander Keith Fuller (US DoD, Defense Information Systems Agency (DISA))
  - Richard Hale (Chief Information Assurance Executive, US DoD, Defense Information Systems Agency (DISA))
  - Carl Landwehr (National Science Foundation (NSF))
  - John Todd (National Communications System (NCS))
  - David Nolan (National Communications System (NCS))
  - Richard Clarke (National Security Council (NSC), Chair of the President's Critical Infrastructure Protection Board (PCIPB) (\*), Special Advisor to the President for Cyber Security)
  - Tommy Cabe ((National Security Council (NSC))
  - Howard Schmidt ((National Security Council (NSC), Vice Chair of the President's Critical Infrastructure Protection Board (PCIPB) (\*))
  - Cengiz Alaettinoglu (Packet Design)

(\*) since resigned

A demo of several scenarios was used to illustrate the benefits of having S-BGP protect an AS from accepting bad routes. A highlight was an attack on BGP routing causing the rerouting of a web browser in a non-S-BGP AS from the correct server to a bad server, followed by a scenario in which a web browser in an S-BGP AS went to the correct server in spite of an attack on BGP routing.

## **2.5 Papers, Internet Drafts, Reports**

- ❖ BBN prepared a paper that was accepted by the DARPA Information Survivability Conference and Exposition (DISCEX II) June 2001. The paper was titled "A Public Key Infrastructure for the Secure Border Gateway Protocol (S-BGP)." It expanded upon previous descriptions of the S-BGP PKI and describes how the initial design has changed.
- ❖ We updated the S-BGP specification to reflect our experience in developing the reference implementation and test system from the prototype.
- ❖ On 2/22/02, we submitted a revised version of the Internet Draft that describes X.509 certificate extensions for IP addresses and AS numbers -- "X.509 Extensions for IP Addresses and AS Identifiers." This draft provides a generic rather than an S-BGP-specific description of the extensions so that other applications besides S-BGP can make use of them. (APNIC staff had expressed interest in having a certificate mechanism for binding IP addresses and AS numbers to the organizations authorized to use/advertise them.) The third S-BGP certificate extension (router IDs) is specific to S-BGP and will be covered in the S-BGP certificate profile document. We revised this Internet Draft during October-December 2003. At the end of December, 2003, the document was under review by the IESG, in preparation for approval as an standards track RFC.
- ❖ In September 2003, Dr. Kent prepared a high level overview paper on the S-BGP project, "Securing the Border Gateway Protocol" The paper was published in The Internet Protocol Journal, vol 6, 3, pp. 2-14.<sup>2</sup>
- ❖ BBN prepared an update to the JSAC paper ("Secure Border Gateway Protocol (S-BGP)", a description of the S-BGP architecture) that was published several years ago. The paper was presented at the 7th IFIP TC-6 & TC-11 Conference on Communications and Multimedia Security (Oct 2-3, 2003), and is reproduced in the conference proceedings. This paper describes the changes that have been made to S-BGP, addresses performance and operational issues that have been raised since the JSAC paper, and assesses related work and competing approaches.<sup>3</sup>

---

<sup>2</sup> This work was done under another contract but was relevant to this project and thus is mentioned here.

<sup>3</sup> This work was done under another contract but was relevant to this project and thus is mentioned here.

- ❖ BBN collected BGP data during October-December of 2003 and analyzed some of it in order to update our assessment of the CPU and memory load imposed by S-BGP (see Final Report, Section 5).

## **2.6 Meetings and Briefings**

- ❖ We attended the DARPA FTN PI meeting in Colorado Springs on 7/30 to 8/2/01. We gave a presentation on this effort and demonstrated initial versions of some of the tools developed to support the S-BGP PKI.
- ❖ In November 2001, we briefed Steve Blumenthal (CTO of Genuity) and several of his key staff on S-BGP.
- ❖ In December 2001, we briefed Peter Zarrella (advisor/deputy to Dawn Meyerriecks (CTO/Chief Technologist of DISA) on S-BGP.
- ❖ In December 2001, we briefed Richard Clarke (special advisor to the President on cyber security) on S-BGP. This resulted in an invitation to brief high-level executives from ISPs, etc in January.
- ❖ We attended the DARPA FTN PI meeting on January 15-18, 2002. We presented a briefing on S-BGP.
- ❖ On 1/30/02, Steve Kent briefed S-BGP to ISPs, router vendors, and a few other companies plus representatives from various government groups such as NCS. The briefing was arranged by Richard Clarke, special advisor to the President on cyber security, to discuss Internet security issues.
- ❖ On 5/9/02, Charlie Lynn attended a meeting convened by the President's Critical Infrastructure Protection Board to discuss how government and the ISP industry can work together on information infrastructure security. This was a meeting of seven ISP/router vendor Working Groups who had been working to identify options to better secure the Internet. In particular, one of the groups was assigned to address DNS and BGP security issues.
- ❖ We attended the FTN PI meeting in San Antonio (January 27-30, 2003), where we gave a status report and demonstrated the system. At each DARPA FTN/DC PI meeting, Doug Maughan selects a few projects for special recognition for the quality and impact of their work. At this meeting, he selected the S-BGP project for "Excellence in Industrial Research".
- ❖ On 6/20/03, Steve Kent gave an S-BGP briefing at NSA to Craig Harber (Chief of NSA's GIG/TC SPO), Chris Kubic (Technical Director of NSA V4), and Bruce Caulkins (Deputy Chief of NSA's GIG/TC SPO). They have expressed interest in using S-BGP in the GIG/TC. In particular, the current architecture calls for multiple

autonomous systems (many of which will be tactical) carrying IP traffic over all sensitivity levels (protected via HAIPE's where necessary) and interconnected via BGP. In this environment, S-BGP seems essential to address critical security vulnerabilities with regard to network availability, traffic analysis, etc.<sup>4</sup>

### **3 Papers, Presentations, Specifications, and Code**

The S-BGP project has produced the following papers, presentations, specifications, and code. They can be found at [www.ir.bbn.com/projects/s-bgp](http://www.ir.bbn.com/projects/s-bgp), which also has copies of documents and presentations from earlier phases of this project.

#### **3.1 Papers**

- ❖ "Securing the Border Gateway Protocol," The Internet Protocol Journal, vol 6, 3, pp. 2-14, September, 2003.<sup>5</sup>
- ❖ "Securing the Border Gateway Protocol: A Status Update," Seventh IFIP TC-6 TC-11 Conference on Communications and Multimedia Security, Springer-Verlag, Lecture Notes in Computer Science, pp. 40-53, October, 2003.<sup>6</sup>

#### **3.2 Presentations**

- ❖ Jan 02 "Securing the Border Gateway Protocol (S-BGP)" a briefing for Richard Clarke's ISP and Router Vendor Workshop
- ❖ Oct 02 Oregon Workshop Meeting Notes
- ❖ Jan 03 DC Workshop Slides illustrating S-BGP router demonstration and NOC Tools

#### **3.4 Specifications**

- ❖ Jul 03 S-BGP Protocol Specification
- ❖ Sep 03 X.509 Extensions for IP Addresses and AS Identifiers 02

#### **3.5 Source Code**

- ❖ Prototype S-BGP source code based on MRT and supporting infrastructure components
- ❖ NOC Tools to manage certificates, CRLs, and Address Attestations
- ❖ Open Source CMS Certification Authority
- ❖ S-BGP Repository

---

<sup>4</sup> This work was done under another contract but was relevant to this project and thus is mentioned here.

<sup>5</sup> This work was done under another contract but was relevant to this project and thus is mentioned here.

<sup>6</sup> This work was done under another contract but was relevant to this project and thus is mentioned here.

## 4 Lessons Learned

- ❖ At least some Cisco routers truncated unknown path attributes at 256 bytes, although a fix for this problem was later made available by Cisco. The S-BGP code was augmented with a per-peer option to remove the Attestation path attribute when sending an UPDATE to accommodate this limitation, consistent with the notion of incremental deployment of S-BGP.
- ❖ The processing needed to verify that each prefix being advertised is covered by each Registration Authority (RA) along the path can require  $O(N^2)$  work per RA. The S-BGP code should sort the prefixes in the Network Layer Reachability Information (NLRI) when generating an RA (which, recursively, translates to the first S-BGP speaker, and aggregation points) so that the work expended by routers receiving UPDATES is linear. This may also apply to AS SETs, but this needs to be explored further as simple appending may be a better solution. (Appending raises the issue of what to do when an AS appears in more than one of the routes being aggregated.)
- ❖ Some operating systems cannot dynamically grow the execution stack. Algorithms that use recursion over the RAs in a path may cause problems; looping (or tail recursion) would be a better technique. Recursion over a certification hierarchy may also be a problem.
- ❖ The large number of public keys needed to validate the full range of UPDATES that a router will encounter (tens of thousands) can significantly reduce the advertised performance of some digital signature hardware.
- ❖ Some digital signature hardware cannot achieve the advertised transaction rates unless a processing pipeline (work queue) is kept full. The flow of control to keep a work queue full may require significant changes from the flow used by BGP. Keeping the work queue full might require many messages to be processed in parallel, increasing the amount of memory needed to hold the in-progress messages -- a typical time-space tradeoff.
- ❖ The information to be signed is often fragmented. Either the information needs to be marshaled, e.g., for some hardware hash implementations, or the hash function needs to be capable of accepting the information in chunks. A software hash might be more efficient than one implemented in hardware in some cases.
- ❖ The original design for the NOC tools had the crypto officer request that the CA sign the certificates requested by an operator -- the operator couldn't do it himself/herself. So whenever the operator created a certificate (e.g., certificates for routers, customers, etc.), the crypto officer had to be involved. The crypto officer does not need to be in the loop for operations for which an operator is already authorized.
- ❖ The NOC tools currently assume that there's a single trust anchor certificate and a single path from the trust anchor certificate to the operator certificate. When one goes

to a level below that of an ISP, the IP addresses come from the ISP but the AS number(s) come from the Regional Internet Registries (RIR). The ISP operator needs 2 certificates -- one for authorization for AS numbers, one for authorization for IP addresses. The NOC tools need to be modified to handle two certificates and figure out which to use.

- ❖ The current NOC Tools combine the applications and the database into a single system. The NOC Tools should be re-structured to allow multiple operators to run the applications (clients) on different workstations and connect to the backend databases (server).
- ❖ While many organizations (ISPs, router vendors, government agencies) said that securing the Internet's routing infrastructure was very important, other issues were more urgent and important. Router vendors and ISPs were distracted by the industry recession and the need to stay afloat. Government agencies were distracted by the recession and 9/11. There was minimal demand from customers or the public that ISPs secure Internet routing.
- ❖ S-BGP needs more memory than most currently fielded routers have. And many COTS routers aren't designed to be upgraded with enough memory to support S-BGP.
- ❖ Most routers can handle the steady state processing load of S-BGP, but not the transient load that occurs when a burst of UPDATES occurs, e.g., when rebooting. This may be less of an issue for some systems, e.g., Juniper, which can add CPUs.
- ❖ ISPs and Military Services require robust, user-friendly, well-documented, supported systems, i.e., good COTS products. Accordingly, tech transfer requires turning a prototype into a COTS system. This requires a lot of time and funding, frequently many times more than the amount needed to develop the prototype.
- ❖ During the past several years, other mechanisms for authenticating Internet resource (e.g., IP addresses) allocation have been proposed. This has created confusion and disagreement over how best to address this problem. It would have simplified things if the IETF standards process for the S-BGP extensions draft had been completed sooner and American Registry for Internet Numbers (ARIN) had issued/signed certificates with the S-BGP extensions.
- ❖ We initially tried to use browsers as the front end for the certificate request and issuance system. Neither Netscape nor Internet Explorer had the necessary features.
- ❖ From the October 2002 workshop, we learned:
  - The workshop participants wanted to protect iBGP routes as well as eBGP routes.
  - The workshop participants wanted to have their own local repositories. This could mean one for each of 100-150 major ISPs and another 100 or so

for exchanges. Also, some ISPs wanted to have a second repository for redundancy.

- The workshop participants wanted to have incremental download (from the repository) capabilities.
- To prevent potential hijacking, a prefix-owner needs to be able to specify that more specific prefixes than the prefix listed in the AA are not allowed. Otherwise an adversary can advertise a more specific route in order to divert traffic to itself instead of letting it follow an assumed more desirable (e.g., "shorter") alternate route. The address attestation format was modified to enable a prefix owner to specify the length of the most specific prefix that it is authorizing the subject AS to advertise.

## 5 BGP Statistics Update

Since the original estimates were made of the load and overhead that addition of the Secure BGP countermeasure causes, the Internet has continued to grow at an exponential rate. The Loc-RIB now (January 2004) contains about 150,000 IPv4 address prefixes and 150 IPv6 prefixes. When the original statistics were analyzed, around 1999, there were 1/3 fewer IPv4 prefixes and essentially no IPv6 prefixes.

The method of statistics collection has changed from that used for the earlier analysis. The original data was collected independently from four ISPs that each provided a BGP peering session. The current data is collected by a single site, [www.routeviews.org](http://www.routeviews.org), that peers with about 40 ISPs, of which 9 are advertising IPv6 routes. The site provides a trace of all BGP UPDATES received from the 40 peers as well as a periodic dump of the contents of the resulting Loc-RIB.

Having access to the Loc-RIB information can improve the estimate of the number of prefixes being advertised. When only UPDATE trace information is analyzed, only those prefixes that were either advertised or withdrawn are observed; prefixes that were continuously present are missed. The fact that the previous estimates based on only looking at UPDATE messages gave answers that were consistent with other ways of analyzing BGP data implies that very few of the prefixes are stable; most are continually being advertised and withdrawn. This result is plausible. If a link along the best path to a prefix goes down, then UPDATES will be sent announcing an alternate route. Thus the prefix will be seen in UPDATE messages. Analysis of trace data sheds little light on the reason that there is so much "churn" in the Internet.

The advantage of having UPDATES from all the ISPs merged into a single file is that the averages are computed over a larger sample, and it is easier to observe variations among the peers. During the month of November 2003, there was an average of about 0.59 IPv4 UPDATES per second per peer (35.43 per minute per peer), averaging about 78 bytes in length. For IPv6, there was an average of about 0.04 UPDATES per second per IPv6 peer (2.56 per minute per peer), averaging 83.1 bytes in length. The average length of an UPDATE has increased. In part this is due to the increased use of the Community path attribute. The Community attribute was present in 54 % of the IPv4 UPDATE messages

and 44 % of the IPv6 UPDATES. If S-BGP is used to protect the Community attribute, the length of the UPDATES will be larger than previously estimated. Several new uses for the Community attribute are being proposed. Use of this attribute should be analyzed to determine whether some or all of the uses can benefit from S-BGP protection mechanisms.

The average AS Path length, including AS pre-pending, during the month was 3.87 (a year ago it was 3.7 and was originally 3.6). The IPv6 average is a little higher: 4.48. By excluding pre-pending, one obtains the number of S-BGP route attestations (RAs) that would be present, and thus the number of digital signatures that would be verified when the UPDATE is received (ignoring optimizations). The IPv4 average is 3.2 and the IPv6 is 4.2. The maximum number of RAs in a single UPDATE that would have appeared in the November data had it been protected by S-BGP was 20. Analysis of earlier data found a similar, though larger number (27). The data showed that the busiest minute in terms of UPDATES received was not the same for the 38 reporting peers. More precisely, if one identified the busiest minute for each router over the 30 day period, the time and day at which each router experienced the busiest minute was not correlated with that of other routers. The difference is too large to be attributed to "dispersion" caused by CPU bound processing..

Only 0.69 % of the IPv4 prefixes were multihomed. Slightly more, 0.84 % of the IPv6 prefixes were multihomed. (Multihoming is defined as observing the same prefix advertised by two or more ASes at the same time.)

Whenever one looks at real network data, one has a few surprises. Analysis of the IPv4 data for November was no exception. The number of IPv4 prefixes found in UPDATE messages was 433854! Looking more closely, 270318 were prefixes, mostly between 63.160 to 63.202, for which the prefix lengths ranged from 11 to 32; there were often multiple prefix lengths for what is essentially the same address block. (Technically, each different length associated with the same base address value results in a distinct prefix, and the routing system treats these prefixes as distinct for purposes of selecting a route, with longer prefixes being accorded preference.) About 22% of the prefixes were of length 24 and another 60% were longer than 24. This seems odd, since many ISPs will not accept prefixes longer than 24! The questionable data was reported by most of the reporting peers (36 of 38). Thus these long prefixes were propagated through a wide portion of the Internet. It is not clear if the propagation is the result of configuration errors in one or more ASes, or if some other phenomenon is to blame. Whether or not the S-BGP mechanisms would have limited propagation of this problem depends on the base cause.